STUTTGART
Internet Exchange

(re-)building an IXP from scratch

Moritz Frenzel, Stuttgart-IX

who am I?

Moritz Frenzel

- Senior Network Architect at ANEXIA
- Volunteer at Stuttgart-IX
- based in Munich, Germany
- previously worked for
  - a CDN
  - a Datacenter
  - an ISP
- Board Member at DENOG
- @momorientes on the internet

Quick Disclaimer

STUTTGART Internet Exchange

My work at Stuttgart-IX is voluntary and none of my current or previous employers or DENOG e.V. encouraged (or discouraged) it in any way.

Stuttgart-IX

STUTTGART Internet Exchange

- 20 Peers, 109G connected capacity
- ~5.6G Peak Traffic
- 4 Locations
  - 1 more to be connected in 2020
- Legally Operated by ISP Service e.G.
- Additional help from 6 volunteers

Preface

- Stuttgart-IX was founded in 2005 by André Scholz and Winfried Haug
  - Winfried later stepped away from the project
- The IXP was operated in good faith until 2015 when André passed away way too early
  - knowledge and access to/about the infrastructure was lost
- A group of volunteers kept the lights on, but updates etc. never occured
- Stuttgart-IX never had sent an invoice, eventhough a pricing had been agreed
- Over all this time the IXP had an amazing community which met for beerings regularly
- In 2019 we asked ourselves and the community how we should move forward

Moving forward

STUTTGART Internet Exchange

Option 1

hand over the peers to one of the more established IXP operators looking to expand into a new Metro

Option 2

rebuild everything, start invoicing and stay independent. Improve setup once money accumulates

Option 3

find someone local who wants to make it their business and hand it over

Option 4

leave it as is and hope it doesn’t break

Moving forward

STUTTGART Internet Exchange

Option 1

hand over the peers to one of the more established IXP operators looking to expand into a new Metro

Option 2

rebuild everything, start invoicing and stay independent. Improve setup once money accumulates

Option 3

find someone local who wants to make it their business and hand it over

Option 4

leave it as is ar doesn't break

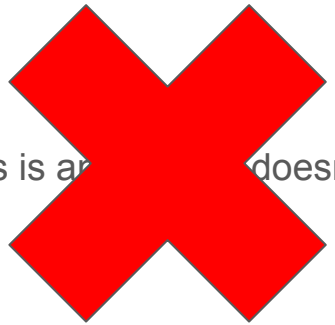

Moving forward

STUTTGART Internet Exchange

Option 1

hand over the peer… …the more established IXP or… …ing to expand into a new Metro

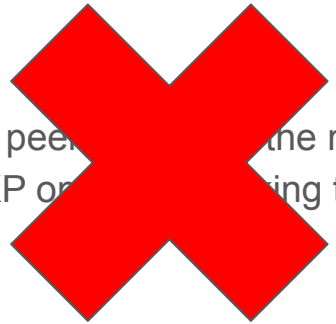

Option 2

rebuild everything, start invoicing and stay independent. Improve setup once money accumulates

Option 3

find someone local who wants to make it their business and hand it over

Option 4

leave it as is a… …doesn't break

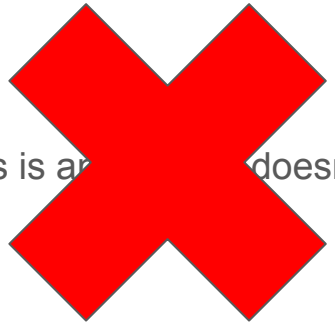

STUTTGART Internet Exchange

Moving forward

Option 1

hand over the peer[illegible] the more established IXP or[illegible]ing to expand into a new Metro

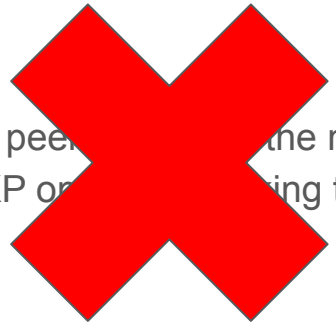

Option 2

rebuild everything, start invoicing and stay independent. Improve setup once money accumulates

Option 3

find someone loc[illegible] make it their business and hand

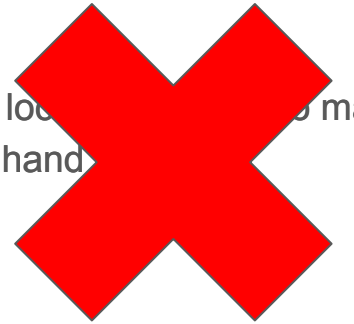

Option 4

leave it as is ar[illegible] doesn't break

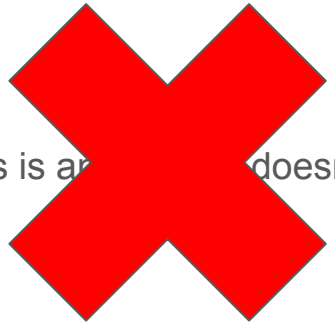

STUTTGART
Internet Exchange

Let’s start from scratch!

STUTTGART Internet Exchange

We had to leave behind

- A Force10 E300 (EOL/EOS since ages, Backplane errors, …)
- Our complete Server infrastructure
  - severely outdated
  - undocumented
  - no access to some parts
- Our Route Servers
  - no access
  - insecure

STUTTGART Internet Exchange

What could we take with us?

- A great community of 18 peers
  - this is probably the hardest to get if you're starting from scratch
- Our Rack and cabling towards our peers
  - completely undocumented
- IPv4/IPv6 Address Space for the Peering LAN
  - If you don't have your own: RIRs have reserved IP space for IXPs
- Our established Domains

STUTTGART Internet Exchange

Goals

- Provide a stable platform for peers to exchange traffic
- Promote local peering
  - get CDN caches etc. on the IXP
- Ensure Route Servers are secure & reliable
- Be MANRS compliant

STUTTGART
Internet Exchange

Everything starts with a switch

STUTTGART Internet Exchange

Switch

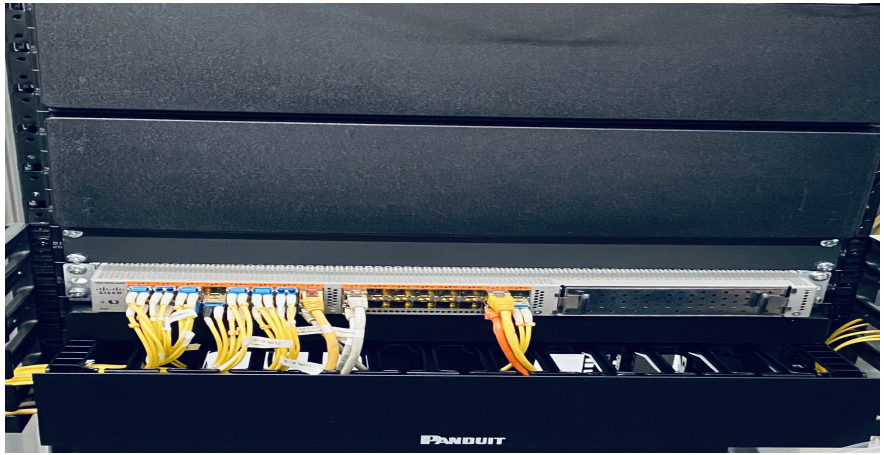

- We had a Cisco Nexus 5548UP and spares available
- This is **NOT** a recommendation
  - Automation is painful
  - Not enough L2 features to securely filter the peering LAN
  - no s/j/netflow only SNMP
- We don’t know what our next switch will be, but we will use the EURO-IX ixp-wishlist to evaluate it

https://www.euro-ix.net/media/filer_public/0a/5b/0a5b4a4e-e032-41f8-b0f7-43c1375c5442/ixp-wishlist.pdf

Switchport Configuration

- Try to protect the peering LAN from all unwanted protocols (CDP, LLDP, …)
- On turnup place peers in a quarantine VLAN and dump their port to ensure their side is configured correctly and is not sending unwanted protocols

```
interface Ethernet1/1
description type=peer, member=BelWue, speed=10000
  no lldp transmit
  no lldp receive
  no cdp enable
  switchport mode trunk
  switchport trunk allowed vlan 100
  spanning-tree port type edge
  spanning-tree bpduguard enable
  logging event port link-status
  logging event port trunk-status
  storm-control broadcast level 1.00
  storm-control multicast level 1.00
  mac port access-group belwue
```

STUTTGART Internet Exchange

Servers

- We acquired 2 new servers to host VMs
  - AMD EPYC 3251 8-Core Processor
  - 64GB RAM
  - 1T SSD
- We used proxmox-ve to spin up new VMs
  - kvm + libvirt would’ve been enough
- Debian Stable on all VMs
- Salt Stack as configuration management

STUTTGART Internet Exchange

IXP Manager

- Manage peers, contact details and users
- Single solution to document everything (switches, ports, IPs, cabling, …)
- Provide statistics and insights to peers (MRTG, smokeping, sflow)
- Renders configuration for route servers
  - Also queries the IRRDBs for prefixes/ASNs and filters accordingly
- Offers APIs for almost everything
  - we use it to automate our switch and DNS infrastructure

STUTTGART Internet Exchange

Route Servers (1VM per physical server)

- birdv2 config generated every 30min by IXP Manager
- routinator on each route server VM for RPKI validation
- birdwatcher provides JSON API towards looking glass
- bird_exporter for monitoring with prometheus

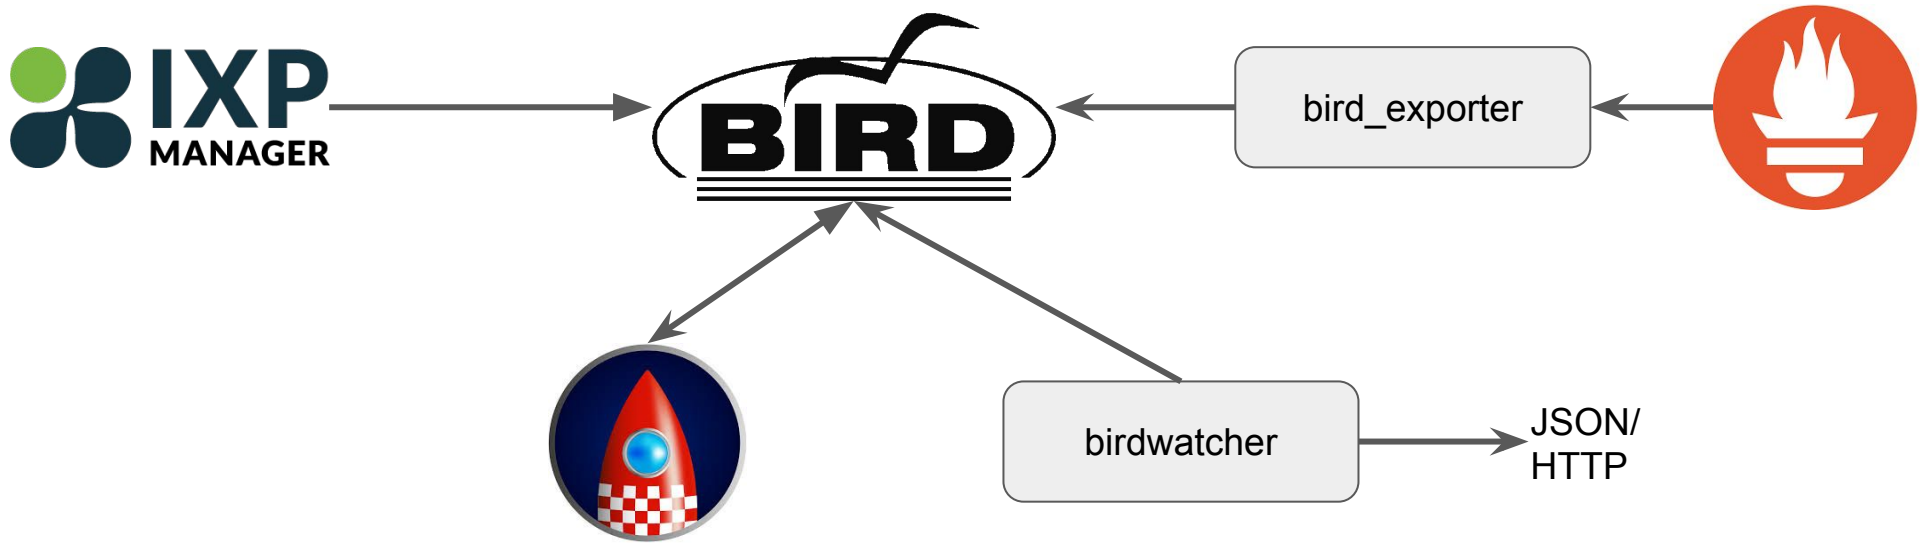

STUTTGART Internet Exchange

Looking Glass

- alice-lg on separate internet facing VM
- queries birdwatcher JSON API every 5m
- we preferred the UI/UX over the IXP Manager looking glass
- great in assisting users to investigate why their prefixes are rejected
- offers a JSON API

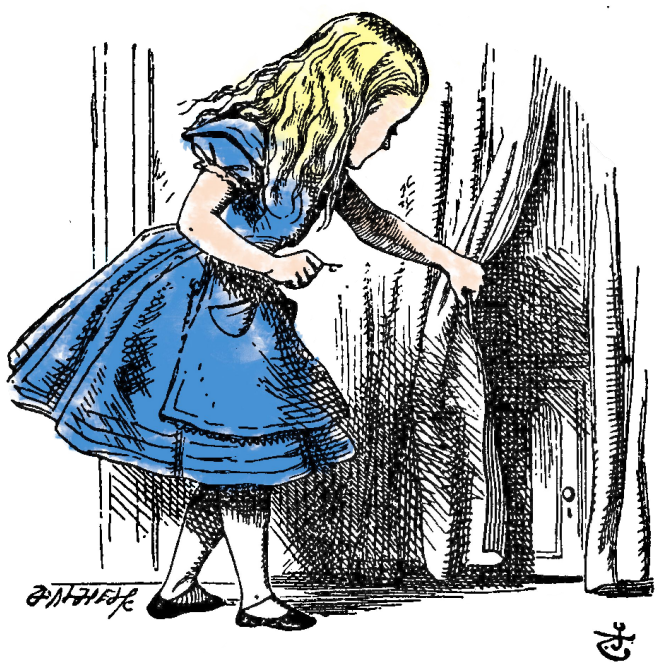

STUTTGART Internet Exchange

Monitoring

- IXPManager integrates with nagios, use it if you have no preference
- We however prefer prometheus, so we use:
  - node_exporter: for server metrics (RAM, CPU, HDD,...)
  - bird_exporter: for statistics on the route servers
  - blackbox_exporter: for icmp/http reachability and SSL certificates
  - cisco_exporter: to monitor our switch
  - nginx_exporter: for our webservers
  - smokeping_prober: as a stand in for smokeping
  - routinator metrics endpoint: to monitor routinator
- Grafana for fancy dashboards

STUTTGART Internet Exchange

Website

Disclaimer: Our website isn’t perfect but here is what I’d expect from an IXPs website:

- Pricing
- Connected peers (maybe even as IX-F JSON Scheme)
- Traffic graphs
- Route server communities & looking glass
- English version

STUTTGART Internet Exchange

PeeringDB

- Create and maintain a PeeringDB record
  - for your IXP
  - for your Route Server ASN
  - for your infrastructure ASN (if you have one)
- encourage Peers to do the same
- Quite a few networks set up peering based on PeeringDB so make sure everything is up to date

DNS (optional)

- we set up PowerDNS with a MySQL backend
  - forward zones are managed by hand
  - reverse zones are managed by IXP Manager
    - peers can request their own PTRs for their peering LAN IPs

STUTTGART
Internet Exchange

Conclusion

STUTTGART
Internet Exchange

= MVP

Education is important

- Especially smaller networks have never heard of routing security and are happy if their BGP setup just works
- We held a 4h long workshop with ~20 participants
  - We removed 30+ invalid prefix announcements from the DFZ
  - Many of them have never heard of RPKI or even IRRDB filtering
  - Some hadn’t even set up max-prefix
- There are great resources available for your workshop:
  - https://www.manrs.org/
  - https://www.euro-ix.net/en/
  - http://bgpfilterguide.nlnog.net/
  - and many more

STUTTGART Internet Exchange

Conclusion

- Building an IXP is fairly straight forward
- We have a great landscape of open source software we can rely on
- Documentation is great and the community is very helpful
- Think hard before buying a switch (it’s the most inflexible part of your IXP)
- We need more pre-packaged software and/or repositories
  - I refuse to curl | sh stuff to install a tool
  - Setting up a rust and golang build-chain requires time & knowledge
- Route Server Security comes at almost no cost (IRRDB filtering + RPKI)

Moving Forward

- To our peers CDNs and Content Providers are important
  - Currently they reach them via transit
  - Waves to Frankfurt + Rack + Port at a big IXP are too expensive
- We've offered them free colo + power + ports + connectivity for cache fill
  - We'll receive a cache by a CDN (50G)
  - 2 politely declined
    - A content provider requires more traffic for a cache (~10G)
    - A CDN doesn't have a standalone solution suited for IXPs
  - 7 never answered
- We've set up a AS112 node
- We're looking to host root-name servers
- Peers are also interested in Cloud & NTP Servers

STUTTGART
Internet Exchange

Questions?

ops@stuttgart-ix.de
mail@moritzfrenzel.de